DRAFT

High-Speed Interconnect for the Space Interferometry Mission[1,2]

W. John Walker
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109
818-354-3260
wjwalker@jpl.nasa.gov

Abstract—This paper describes a high-speed avionics interconnect for a space-borne stellar interferometer. To meet its objectives, the interferometer must make extremely precise measurements in the presence of disturbances due to vibration and thermal effects. A large number of actuators and sensors separated by distances up to several meters and interconnected by interacting digital control loops must operate synchronously at sample rates up to 4kHz. The number of devices, their physical separation, the high rate at which they must be serviced, the degree to which measurements and actuations must be coordinated in time and the physical distances spanned place unusually stringent requirements on the communication fabric that connects the sensors and actuators with the flight computers which implement the control laws. Additional challenges are posed by the need to withstand the rigors of launch and several years of unattended operation in space. The project is in formulation phase and many requirements are subject to change, but the current baseline design has enough definition to allow the essential characteristics of a suitable interconnect to be discerned. The key requirements will be described and the applicability of two possible approaches to meeting them will be briefly surveyed.

TABLE OF CONTENTS

1. THE SPACE INTERFEROMETRY MISSION (SIM)

In 2008, NASA's Space Interferometry Mission (SIM) will place in earth-trailing orbit a stellar interferometer that is capable of measuring the positions of stars and other astronomical objects with an accuracy of 4 microarcseconds. Figure 1 illustrates the SIM spacecraft. The ability to perform astrometric measurements with this unprecedented accuracy will enable astronomers to investigate many hitherto unobservable phenomena. Of particular interest is the potential for detecting earth-like planets in orbit around a distant star by measuring extremely small perturbations in the star's position due to the gravitational pull of the planet.

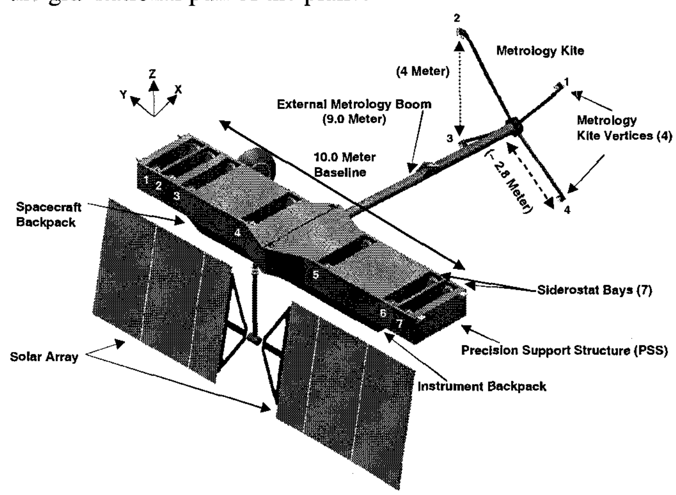

Figure 1 SIM Spacecraft

2. STELLAR INTERFEROMETRY

The essential function of a stellar interferometer is to measure the angle between a "baseline" drawn between

[1] 0-7803-6599-2/01/$10.00 © 2001 IEEE
[2] Updated Oct 23, 2000

fiducial points on two spatially separated siderostats and the incoming wavefront from a target star.

The basic operating principle is that white light fringes will form when the optical pathlengths traversed by parallel starlight beams from the two siderostats are nearly equal, and a bright central fringe will form at the point where the pathlengths traversed are exactly equal. Figure 2 illustrates the geometry.

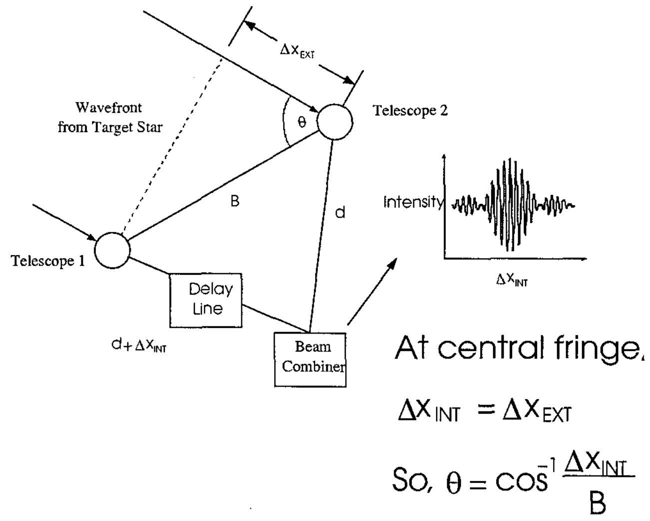

Figure 2 Geometry for Astrometric Measurement

The distance traveled by starlight incident on Siderostat 2 will be greater than that which is incident on Siderostat 1 by an amount equal to the *external pathlength difference* Δx_{Ext}.

If the external pathlength difference can be offset by using an optical delay line to increase the pathlength in the shorter arm until white light fringes form, then the *internal pathlength difference* Δx_{Int} between the optical pathlengths traversed by starlight in the two arms of the interferometer will be equal to the external pathlength difference Δx_{Ext} at the location of the bright central fringe.

If the internal pathlength difference and the length of the baseline B can be measured, then the desired angle can be calculated as indicated in Figure 2.

3. The SIM Instrument

A single interferometer can only make astrometric measurements with respect to a single baseline embedded in the spacecraft.

To meet SIM's objectives, it is necessary to establish the orientation of the baseline in inertial space. While the spacecraft will have a conventional attitude determination system that can provide an initial estimate of baseline orientation, its accuracy will be orders of magnitude too coarse for astrometric purposes.

To obtain the required fine attitude knowledge, SIM uses two "guide" interferometers whose baselines are parallel to that of the "science" interferometer to track small changes in the orientation of the common baseline with respect to bright guide stars whose inertial positions are known.

Real time knowledge of baseline orientation obtained by tracking the central fringes of two guide stars is used to generate *pathlength feedforward* commands for the science delay line which is used to stabilize the phase of the target star's fringe pattern.

To allow the fringes for a dim science star to be imaged with an acceptable signal to noise ratio by a detector with finite sensitivity, this phase stability must be maintained with sufficient accuracy for several minutes.

Equally important is the task of stabilizing the angular alignment of the beams in the two arms of the interferometer with respect to the line of sight to a science star that is too dim to be used as a pointing reference.

This is accomplished by using pointing information from the four guide interferometer siderostats to develop *angle feedforward* commands for the actuators which control the alignment of the science beams.

4. Sensors and Actuators

To make the required astrometric measurements, it is necessary to establish precise real time control over the optical pathlengths between optical elements located at various points on the spacecraft structure. The instrument incorporates numerous sensors and actuators to accomplish this, and they will now be briefly described.

Starlight Sensors & Actuators

Seven *Siderostats* are located at intervals along the 10 meter length of a *Precision Support Structure* (see Figure 1). Each siderostat is a 33cm flat mirror that is articulated in two axes by *Gimbal Actuators* to facilitate acquisition of stars within a 15° field of regard.

Gimbal Encoders measure the orientation of the gimbals with respect to the spacecraft structure.

A telescope compresses the starlight collected by each siderostat into a 3cm beam.

Misalignment of the incoming starlight from the correct boresight as it enters the telescope can result in it being translated in an unpredictable way across the surface of subsequent elements in the optical train. This *beamwalk* effect is a significant source of astrometric error. To control it, a portion of the incoming starlight beam is directed into a *Siderostat Camera* which measures its angular deviation from the desired boresight.

Three interferometer s (two guide plus one science) can be formed by combining starlight from any two of the seven siderostats in any three of four *Beam Combiners*.

Since mission objectives call for only three interferometers, only six siderostats and three combiners are strictly necessary. The seventh siderostat and the fourth combiner are spares which allow the mission to proceed in the event of a malfunction that makes one siderostat or one combiner unusable.

In each interferometer, the two beams pass through an optical train which includes a two axis *Fast Steering Mirror* to maintain their angular alignment, and an *Optical Delay Line* which allows the external pathlength difference to be offset.

Each starlight beam arrives at a Beam Combiner and is divided equally between two optical trains: a fringe tracker and an angle tracker.

In the fringe tracker, the beams from the two arms are combined and a *Fringe Tracker Camera* senses the distance of the central fringe from a point that is exactly equidistant from the left and right siderostat fiducials when the delay line is in its null position.

In the angle tracker, a telescope focuses the angle tracker beams from the two arms onto two distinct spots on the detector of an *Angle Tracker Camera*. The lateral and transverse deviation of each spot from a fixed reference point on the detector is a measure of the corresponding beam's angular deviation from a path which would cause the corresponding fringe tracker beam to fall on a corresponding datum in the Fringe Tracker Camera's detector.

Metrology Sensors & Actuators

As noted previously, it is necessary to measure the difference Δx_{INT} between the optical pathlengths traversed by the starlight in the two arms of the interferometer, and also the length of the baseline vector B between the two siderostat fiducials. These measurements must be accurate to within a small fraction of the wavelength of visible light and are made using laser metrology.

A *Metrology Gauge* accepts laser light with two orthogonal polarizations that differ slightly in frequency, and causes one polarization to traverse the unknown distance while the other one does not. Because of the pathlength difference, there will be a phase shift between a heterodyne signal derived by interfering the two polarizations upstream of the unknown path and a second heterodyne obtained by interfering them downstream of it. Measuring this phase shift electronically gives a relative measure of the unknown distance.

Each of the four Beam Combiners needs 2 *Internal Metrology Gauges* to measure the distance traversed by the starlight beams as they travel from the starlight fiducials to the combination point. Great care is taken to ensure that the metrology and starlight beams are co-boresighted and have pathlengths that are exactly equal.

This close alignment is exploited to assist with fine pointing of the science interferometer at dim target stars. In each siderostat, part of the internal metrology beam is split off and used to illuminate a *Quad Cell*. Because the starlight and metrology beams are well-aligned, the lateral and transverse displacement of the metrology beam from the center of the quad cell is a measure of small high frequency misalignments that arise due to flexing of the structure on which the siderostat is mounted.

Distance measurements from 34 *External Metrology Gauges* are used to establish the lengths and relative orientations of the baseline vectors between paired siderostats by a triangulation process. Associated with each siderostat are four gauges which measure the distances from the appropriate siderostat fiducial to the four vertices of a 2 meter square *Metrology Kite* which is located at the end of a 12 meter boom. A further 6 External Metrology Gauges are needed to measure the 4 sides and 2 diagonals of the kite.

For external metrology triangulation, a kite with only three vertices are strictly necessary. The fourth vertex is included for redundancy.

Kite Vertex Translation– To compensate for thermal deformations in the structure, each of the kite vertices is articulated by three *Kite Vertex Actuators* which allow it be translated in three orthogonal directions.

5. SIM AS A CONTROL SYSTEM

The SIM instrument can be regarded as a closed loop control system in which the respective outputs and inputs of various sensors, flight computers and actuators are tied together by the high speed interconnect which is the subject of this paper. See Figure 3.

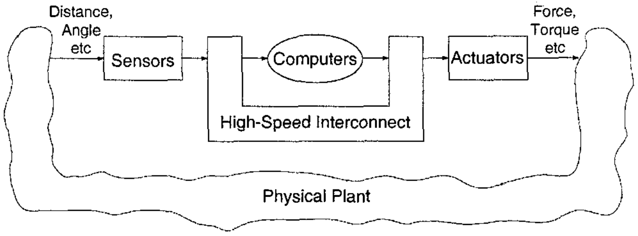

Figure 3 Role of Interconnect in SIM Instrument

It is useful to characterize each sensor as a transducer which converts a physical quantity such as distance, angle or illumination into one or more digital measurements, and each actuator as a transducer which converts digital commands into a physical stimulus such as force or torque.

The role of the flight computers is to transform digital measurements into digital commands by executing real-time control algorithms which are designed to ensure that certain angles and distances are held at prescribed values within prescribed error bounds that are ultimately traceable to the mission requirements (such as 4 uas astrometry).

Each of the three interferometers depends on three major real time control algorithms: Pathlength Control, Wide Angle Pointing, and Narrow Angle Pointing.

The *Pathlength Control (PLC) Task* uses measurements from the Fringe Tracker Camera and two Internal Metrology Gauges to equalize the internal and external pathlength differences by commanding the Optical Delay Line.

The *Wide Angle Pointing (WAP) Task* uses measurements from two Siderostat Cameras and the associated Gimbal Encoders to keep the incoming starlight beams aligned with the appropriate boresight as they enter the compressor telescope by commanding the Gimbal Actuators.

The *Narrow Angle Pointing (NAP) Task* uses spot position measurements from the Angle Tracking Camera to ensure that the starlight beams are parallel when they enter the beam Combiner by commanding the Fast Steering Mirror.

The science interferometer depends on three additional tasks: Baseline Estimation, Pathlength Feedforward and Angle Feedforward.

The *Baseline Estimation (BLE) Task* uses distance measurements from the External Metrology Gauges and pathlength estimates from the Pathlength Control Tasks in the two guide interferometers to estimate the baseline vector in inertial coordinates.

The *Pathlength Feedforward (PLFF) Task* uses the baseline estimate to generate a delay line command for the for the Science Interferometer's Pathlength Controller.

The *Angle Feedforward (AFF) Task* uses Gimbal Actuator commands from the two guide interferometers to generate a feedforward input for the Science Interferometer's Pointing Controller.

Figure 4 illustrates the relationship between the real time control tasks just described and the sensors and actuators described in Section 3.

6. REQUIREMENTS

Fault Containment Regions

The SIM instrument is required to be single-failure tolerant. In other words, it must be capable of completing its mission in the presence of any single failure (from a defined set).

Since the response to this requirement involves redundancy, it is necessary to inhibit the propagation of faults between redundant elements by partitioning the hardware into *Fault Containment Regions.*

A Fault Containment Region (FCR) is an assembly or group of assemblies which is designed so that no fault (from a defined set) occurring within it can cause irreversible damage or permanently degraded performance in any external assembly*.

* Partitioning the hardware in this way leads to a more robust system because it builds in low-level barriers to the propagation of hardware failures, and provides a solid basis for strategies which detect, isolate and recover failures in flight.

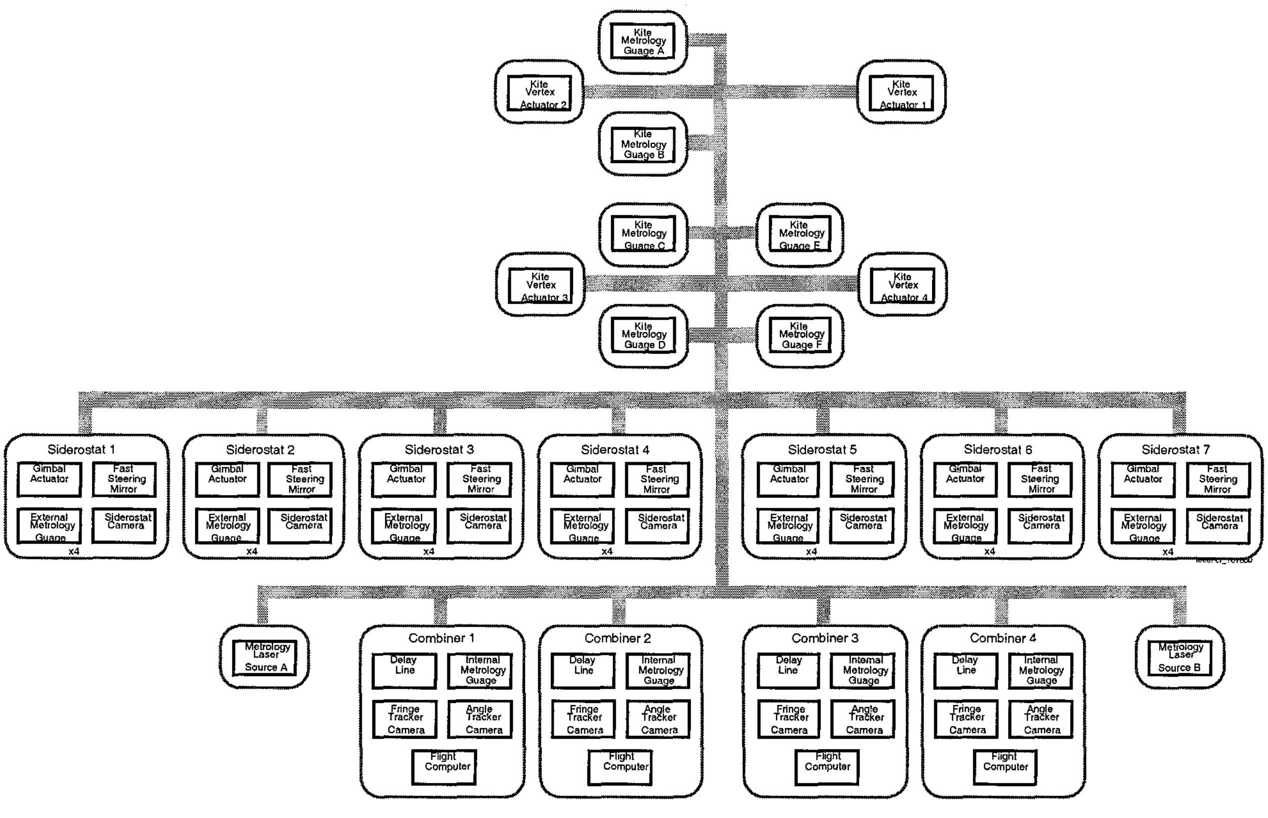

Figure 5 Instrument Fault Containment Regions

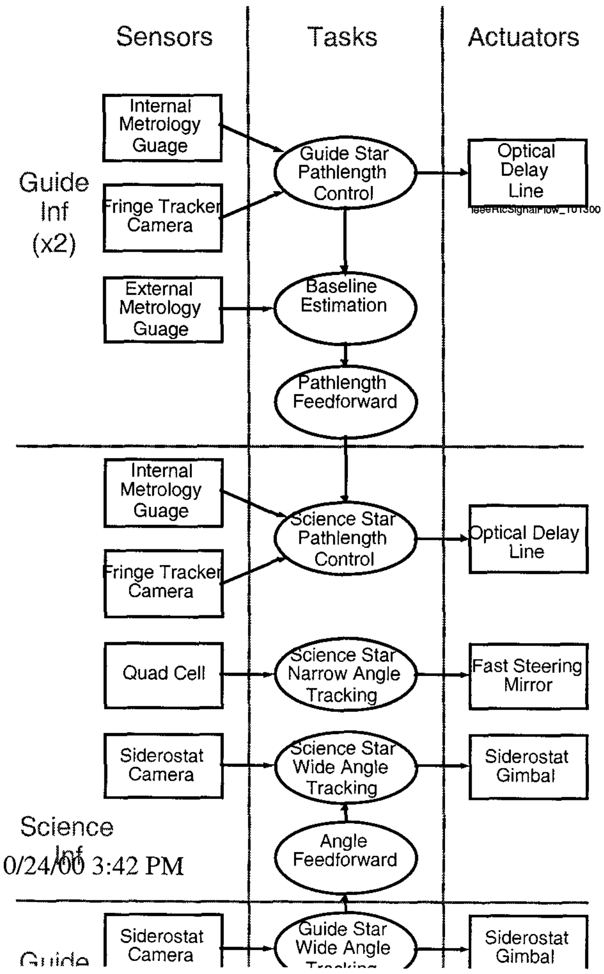

Figure 4 Sensors, Tasks & Actuators

There are 23 Fault Containment Regions in the SIM instrument and they are illustrated in Figure 5.

The Fault Containment Regions are largely determined by the high-level architecture of the instrument, the configurations in which it is required to operate and the needs of the functions they enclose.

Each Fault Containment Region receives power from the spacecraft through a switch (so that it can be disconnected in the event of a fault) and communicates digitally with other nodes via the high speed interconnect.

Each fault containment region contains power converters, an *I/O Node* (which acts as a bridge to the interconnect) and A/D and D/A circuitry that interfaces analog actuators and sensors to the appropriate I/O node.

From the foregoing discussion it will be clear that the interconnect must be capable of supporting at least 23 nodes, that it must provide at least two independent communication channels between each pair of nodes that need to exchange information, and that its interface with each node must be designed so that known hardware faults cannot propagate from the interconnect to the node or vice versa.

Number and Physical Location of Flight Computers

Given a sufficiently capable interconnect, the architecture in Figure 5 suggests a decentralized processing scheme in which the tasks described in Section 5 are executed by Flight Computers that are distributed among all - or at least many - of the 23 nodes.

However, the demands placed on the interconnect by such an approach would be considerably greater than those which would pertain if its only role was the provision of "traditional" i/o facilities (albeit at high speed) for a smaller number of centralized computers. This suggests that the optimal number of independent Flight Computers significantly less than 23.

On the other hand, the number of independent Flight Computers must certainly be at least two to preclude scenarios in which a single processor failure could lead to mission failure.

Another consideration is the total required processor throughput. The question of how to characterize the throughput needs of an interferometer in terms that can be related to performance measures commonly used in industry is the subject of much on-going work at present, but a very preliminary estimate of the total SIM requirement (for 3 interferometers) is 1500 MIPS.

Since the advertised capability of currently available flight qualifiable CPU's is in the vicinity of 250 MIPS, a configuration with fewer than six CPU's is unlikely to be suitable.

The assignment of CPU's to specific Fault Containment Regions is also influenced by the fact that the ability of a local processor to do useful work is not the same for all nodes. It is maximized in the case of the Combiner nodes because each of them contains both a Fringe Tracking Camera and the corresponding ODL.

A 500 MIPS Flight Computer in each of the four Combiner nodes would therefore be able to perform all the fringe tracker processing for a single interferometer (which accounts for a significant fraction of the total CPU load) without relying on the interconnect for I/O transactions.

One of these Flight Computers would be designated as Master for the purpose of configuring and coordinating the system.

The baselined assignment with one flight computer in each of the four combiner nodes is shown in Figure 5. It represents a compromise between the robustness afforded by a larger number of distributed processors, interconnect complexity, the capabilities of available CPU hardware and the fact that the leverage exerted by a local CPU is not the same for all nodes.

Inter-Node Distances

The maximum inter-node cable run is an important constraint on the interconnect. It is determined by the physical disposition of the sensor/actuator clusters, the architecture of the instrument, and the physical extent of the electronics.

The maximum straight line distance between the nodes is in the order of 10 meters.

The maximum cable run that the interconnect must support depends on whether the A/D and D/A electronics are *colocated* with the associated actuators and sensors or *centralized* at a location that may be several meters away.

Support electronics should ideally be colocated rather than centralized because this simplifies cabling, facilitates decentralized testing and tends to mitigate EMC problems. This would imply a maximum inter-node cable run of about 30 meters (allowing a factor of 3 for routing).

On the other hand, the advantages of colocation must be set against structural and dynamic issues that arise when electronics assemblies with finite mass and volume are attached to slender flexible structures, and thermal issues that arise because the waste heat they generate causes problems for the millikelvin temperature controllers that are required to maintain dimensional stability in the optical train.

With the electronics centralized at a remote location, the maximum inter-node cable run for the interconnect would certainly be less than 30 meters, but not much less than the largest dimension of the total volume occupied by the instrument electronics itself. Given the large complement of sensors and actuators, this volume is quite large and preliminary estimates indicate that the maximum inter-node cable run would probably not be less than 5 meters.

So the maximum inter-node cable run requirement could be reduced from 30 meters to 5 meters by changing from colocated to centralized electronics. This is not a sufficient reduction to have a significant impact on the interconnect requirements.

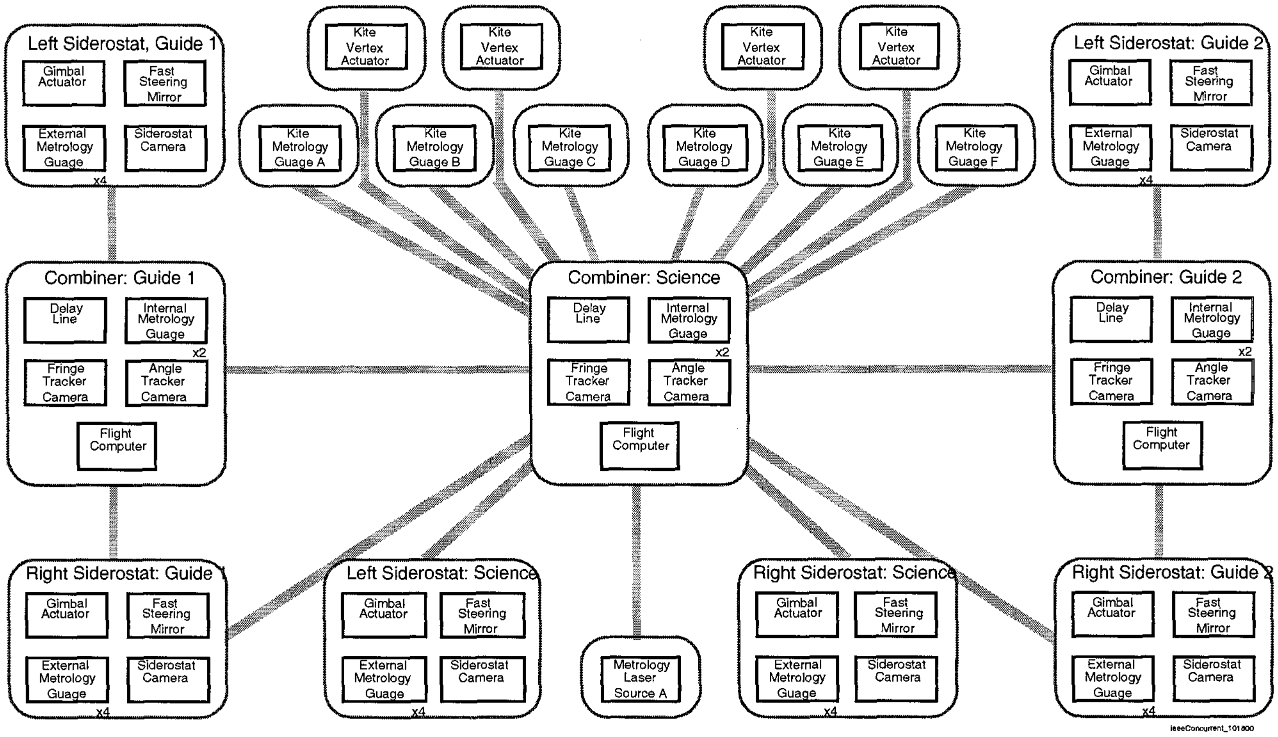

Figure 7 Concurrent Connectivity

If centralizing the electronics had been capable of reducing the maximum cable run to well under 1 meter instead of 5 meters, a number of traditional backplane interconnect solutions such as VME and PCI that can support very high speed over very short distances would become feasible.

As it is, an interconnect that can support a 30m hop is more problematic than one that can support only 5m, but not dramatically so, and 30 meters has therefore been accepted as the driving requirement on maximum inter-node cable run.

Hardware Connectivity

The interconnect must support the formation of 2 guide interferometers and 1 science interferometer by connecting:

- any 6 out of 7 siderostat nodes to any 3 out of 4 combiner nodes.
- any 1 of the 4 combiner nodes (the science node) to 6 kite metrology gauges, 4 kite vertex actuators and 1 laser source
- any combiner node to any of the other 3 combiner nodes (to support the Pathlength and Angle Feedforward tasks).

Figure 6 illustrates the data paths that the interconnect must provide.

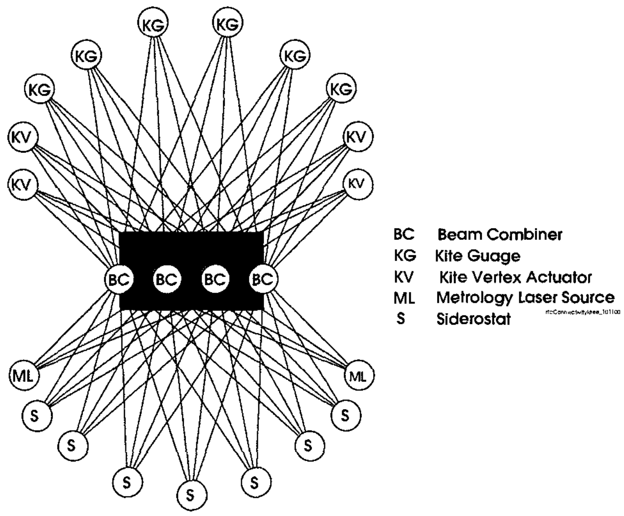

Figure 6 Hardware Connectivity

The presence of a link between two nodes in Figure 6 indicates that the interconnect must support the transmission of data between them. Data needs to be able to flow in both directions, but not necessarily at the same time.

Figure 6 is not meant to imply that the interconnect necessarily has to be implemented using point to point links. In fact, a point to point implementation is unattractive because of the number of links that would be required, the difficulty of enforcing fault containment boundaries and the sensitivity of such a design to small changes in the instrument architecture.

This last point is important because the baseline architecture for the SIM instrument that has been presented here is provisional and may well change significantly before launch. It is therefore important for the interconnect implementation that is chosen to be sufficiently flexible to accommodate the addition or deletion of nodes and/or links from the connectivity illustrated in Figure 6 without major redesign.

Concurrent Connectivity

All the links shown in Figure 6 need to be present to support all the required configurations, but only a subset needs to be active in any given operating mode. This subset is illustrated in Figure 7.

Throughput: A Lower Bound

Table 1 displays preliminary estimates of the amount of data generated or received during each update cycle by each device that contributes significantly to traffic on the interconnect, and the maximum rate at which it must be serviced. Several devices (such as the Kite Vertex Actuators and the Metrology Source) that do not contribute significantly are omitted. Totalizing the data rates as shown leads to a lower bound on the required throughput of about 625kByte/s. While this is not exceptionally challenging in itself, it is high enough to rule out many flight-proven interconnects (such as Mil-Std-1553 and numerous variants thereof) that would otherwise be considered attractive. It should also be emphasized that this figure is just a lower bound: much higher throughputs will need to be sustained over short periods within each update cycle because scheduling and synchronization constraints will inevitably make it impossible to distribute traffic uniformly.

Latency

Any delay that is incurred between the point at which a fresh sensor measurement becomes available and the point at which the appropriate actuator(s) begin to execute commands derived from it will have a detrimental effect on closed loop performance.

In an interferometer, closed loop performance requirements are most conveniently expressed as statistical representations of pathlength error and pointing error. The dynamics that govern the relationship between these errors and I/O latency are complex, but it is possible to bound the required latency by considering the closed loop bandwidths needed to suppress the principal on-orbit disturbances due to reaction wheel imbalance and so on.

Such analysis[1] has shown that among all the control loops that depend on the high speed interconnect, the highest required update rate is 1 kHz and that the maximum total delay that can be tolerated at this update rate is $400\mu s$.

Table 1. Average Throughput

Node	Device	Device Update Rate Hz	Message Length Byte	Quantities Guide 1	Guide 2	Science	Total	Interconnect Throughput kByte/s
Combiner	Pathlength Feedforward	1000	8	0	0	1	1	8
	Angle Feedforward	1000	8	0	0	1	1	8
Siderostat	Sid Ext Met Guage	1000	12	8	8	8	24	288
	Siderostat Camera	200	128	2	2	2	6	153.6
	Gimbal Actuator Fine	200	4	4	4	4	12	9.6
	Gimbal Actuator Coarse	200	2	4	4	4	12	4.8
	Gimbal Encoder	200	4	4	4	4	12	9.6
	Fast Steering Mirror	1000	12	2	2	2	6	72
Kite	Kite Ext Met Guage	1000	12	0	0	6	6	72
							Total	625.6

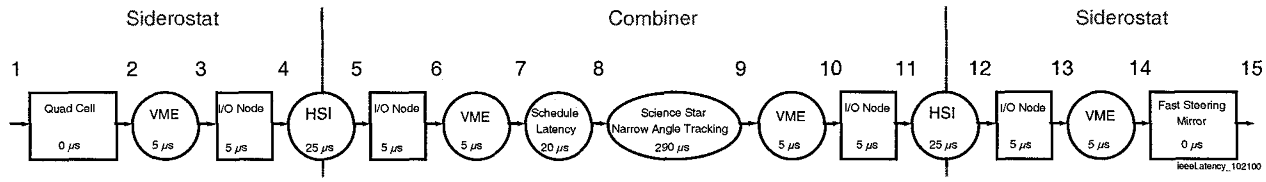

1. Quad Cell makes a measurement.
2. Siderostat I/O Node begins readout of Quad Cell measurement via Siderostat VME.
3. Measurement resident in Siderostat I/O Node memory.
4. Siderostat I/O Node begins transmission of measurement to Combiner I/O Node via HSI.
5. Measurement resident in Combiner I/O Node memory.
6. CPU begins readout of measurement from Combiner I/O Node via Combiner VME.
7. Command resident in CPU memory.
8. CPU executes first instruction in Narrow Angle Tracking task.
9. CPU begins transmission of command to Combiner I/O Node via Combiner VME.
10. Command resident in Combiner I/O Node memory.
11. Combiner I/O Node begins transmission of command to Siderostat I/O Node via HSI.
12. Measurement resident in Siderostat I/O Node memory.
13. Siderostat I/O Node begins transmission of command to Fast Steering Mirror Electronics via VME.
14. Command resident in Fast Steering Mirror Electronics memory.
15. Fast Steering Mirror begins execution of command.

Figure 8

To assess the portion of this total delay that can reasonably be allocated to the high speed interconnect, it is necessary to consider all the incremental delays that contribute to it. The worst case arises in the Narrow Angle Pointing loop for the Science Interferometer because both sensor (Quad Cell) and actuator (Fast Steering Mirror) are located in a siderostat that is remote from the flight computer which is located in a combiner.

As signal flow proceeds from sensor to actuator, the most significant delays that are encountered may be classified as follows:

- delays due to VME Bus communication between sensors/actuators and the I/O Node
- delays that occur within an I/O Node
- delays due to the High Speed Interconnect
- scheduling latency within the Flight Computer
- task execution time within the Flight Computer

Figure 8 illustrates a tentative assignment of delay to each of these items as they occur in the Narrow Angle Pointing Loop for the Science Interferometer.

Synchronization & Timing Jitter

An important characteristic of the SIM instrument is that certain sensor measurements must be synchronized with an uncertainty equivalent to a small fraction of the update period.

This requirement has important implications for the interconnect and a strategy for meeting it is illustrated in Figure 10.

Each of the four Flight Computers has the ability to generate a Timing Reference Signal.

The TRS from whichever one of the flight computers is designated Master is distributed to all the nodes through the physical layer of the interconnect.

Within a remote node, each sensor contains hardware that allows it to initiate a measurement on a designated edge of the TRS, as specified by prior transactions between logic in the local I/O Node and the Master Flight Computer.

If this mechanism is used to set up all the distributed sensors to start taking a measurement on the same edge of the TRS, the variance of the time at which the measurement is valid will be limited primarily to whatever timing skew is acquired by the TRS as it propagates between the various nodes.

The amount of difficulty posed by this requirement obviously depends on the exact value of the required accuracy. This has yet to be determined but it will probably lie between $0.1\mu s$ and $10\mu s$.

Qualifiability

In addition to meeting the performance requirements discussed, the hardware used to implement the interconnect must be capable of being qualified for operation in the flight environment in terms of exposure to appropriate levels of shock, vibration, ionizing radiation, temperature extremes and electromagnetic interference.

In the ideal situation where an off-the-shelf implementation of an existing design that has been previously qualified in a similar environment is available, qualification would not be a major issue.

However, if (as seems likely) no off-the-shelf solution is directly applicable without modification, then modified or newly-developed hardware will need to be qualified for flight, and the cost and schedule risk associated with this activity may be an important basis for discriminating between options.

Testability

The interconnect hardware must provide for test access during integration and test. The effort required to develop suitable support equipment for this purpose is a significant cost that must not be overlooked. In the event that more than one solution meets the basic functional requirements, commercial availability of suitable test gear and support software for a particular candidate would therefore be a significant advantage.

7. Implementation Options

The ideal implementation scenario for the requirements described in Section 6 would be adoption of an existing bus standard that could be used with little or no modification.

It is therefore natural to begin by looking at the many serial avionics buses that have been used extensively in the past. Buses in this category that have been considered include MIL-STD-1553 and many variants thereof, ARINC-429 and ARINC-629. These buses have many of the required connectivity, fault containment and redundancy properties and abundant flight heritage. Unfortunately all of them were ruled out at an early stage because - without radical modification – none has sufficient throughput capacity to approach the lower bound described in Section 6.

Another option that was considered and rejected was the use of a backplane bus such as VME as the sole basis for intercommunication. Such an approach is superficially attractive because VME would certainly have no difficulty supporting the point-to-point throughput and latency requirements, and also because the ground based interferometry programs from which SIM is derived and the testbeds upon which ongoing SIM technology development is proceeding make extensive use of VME. The most fundamental difficulty is that any backplane bus has physical limitations which make high speed incompatible with SIM's long inter-node distances, large node counts and extensive connectivity requirements.

IEEE-STD-1394A, also known as "Firewire" is a commercial standard which is being developed and adapted for use in flight systems at JPL[2]. It is capable of meeting SIM's fault tolerance, connectivity and throughput requirements, but is not able to meet the $25\mu s$ latency requirement or the 30m meter requirement on inter-node distance.

IEEE-STD-1596 (1992), also known as Scaleable Coherent Interface is an standard that was developed as a high throughput, low latency memory sharing interconnect for inter-processor communication in high performance clusters. It provides bus-like services but avoids bottlenecks in traditional bus systems by exclusive use of very fast unidirectional point-to-point links. Its main weaknesses in the SIM context are its lack of flight heritage and the fact that there is no ready-made approach to SIMM's requirement for fault tolerance.

IEEE-STD-1393 (1999), also known as Space Fiber Optic Data Bus is currently under development as the basis for a uniform approach to on-board spacecraft data handling3.

Like SCI, it appears to have the potential for meeting SIMM's performance requirements. Its main advantage over SCI is that it is being developed specifically for operation in the space environment and so qualifiability and fault tolerance have been major issues from the outset.

Table 2 summarizes the most significant properties of the options that have been considered.

Table 2

	Fault Tolerance	Inter-Node Distance	Connectivity	Throughput	Latency	Synchronization	Qualified	Commercial Testgear
MIL-STD-1553	Y	Y	Y	N	N	N	Y	Y
ARINC-429	Y	Y	Y	N	N	N	Y	Y
ARINC-629	Y	Y	Y	N	N	N	Y	Y
MIL-STD-1773	Y	Y	Y	N	N	N	Y	Y
IEEE-STD-1394A	Y	N	Y	Y	N	N	Y	Y
IEEE-STD-1596	N	Y	Y	Y	Y	N	N	N
IEEE-STD-1393	Y	Y	Y	Y	?	N	N	N

8. CURRENT STATUS

As stated earlier, the processes whereby high-level interferometer performance measures flow down to timing and throughput constraints on the interconnect is the subject of much on-going work. Ensuring that the implications of these complex dynamics are adequately reflected in the interconnect requirements is the most important near-term goal.

At present, SCI and SFODB appear to be the most promising implementation options. Commercial SCI hardware is being evaluated as a basis for inter-cage communication in the SIM testbed at JPL and results to date have been encouraging.

ACKNOWLEDGEMENTS

The author wishes to thank the C. Bell, M. Tuszyinski and M. Wette, his colleagues at JPL and D. Schaechter at Lockheed Martin for many helpful discussions during the preparation of this paper. This work was performed for the Jet Propulsion Laboratory, California Institute of Technology, sponsored by the National Aeronautics and Space Administration.

REFERENCES

[1] Schaechter, David
[2] Chau, Savio
[3] IEEE-STD-1393